



IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: Jeffrey S. Kuskin, et al. :
Serial No: 09/662,991 : Art Unit #2136
Filed: 15 September 2000 : Examiner:
Title: KEY CACHING SYSTEM : C. Colin

DECLARATION OF TAO-FEI SAMUEL NG UNDER 37 C.F.R § 1.131

Mail Stop AMENDMENT
Honorable Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

I, Tao-Fei Samuel Ng, do hereby declare as follows:

1. I am one of the co-inventors in the above-referenced Patent Application, and presently hold the position of Senior Logic Designer at Atheros Communications, Inc. (Atheros), the Patent Application's assignee of record.
2. The subject matter of our invention, which I understand to be disclosed and claimed in this Patent Application, was developed to the point of actual test integration at Atheros before July of 2000. By at least as early as May of 2000, the other co-inventors and I had developed simulated implementations of the invention for testing and

engineering refinement purposes.

3. Based on such simulated implementations, we integrated certain versions of our invention into various developmental systems, including implementation in integrated circuit form. We did so with the objective of realizing a reliable and refined implementation that could be marketed by Atheros. Attached as EXHIBIT A is a copy of an excerpted portion of a confidential hardware specification document prepared for internal distribution within Atheros, dated 5 May 2000, for one such developmental system. The document includes descriptive mention of certain aspects of the encryption key-caching provided by our invention as integrated in this particular system.

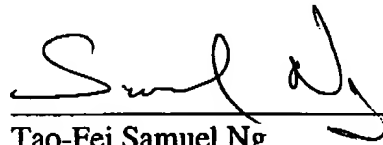
4. Also attached, as EXHIBIT B, is a copy of a brief written summary describing illustrative operation of the encryption key-caching mechanism incorporated into communications equipment in accordance with our invention. The summary, dated "6/21/00," was prepared for confidential internal distribution within Atheros.

5. By at least the latter part of June 2000, we had provided the invention's disclosure information to patent attorneys for Atheros so that they could prepare a suitable patent application on our invention for filing at the U.S. Patent and Trademark Office.

I declare further that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and

further, that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, or any Patent issuing thereon.

Signed this 29th day of September, 2005.



Tao-Fei Samuel Ng



P4DB Version 0.99f

Current change level:
147437

File

//depot/projects/t2/f2/pcu/doc/pcu.txt#1

Main

- Change number -- see the change description
- Revision number -- see diff at selected revision

Download file

Type: text

Line	Ch.	Rev
	1190	1

| Time Bases

| In order to support various real-time related functions eg. TSF, timers,

5: | Inter-frame spacings, NAV, random back-off and ACK/CTS time-out, PCU maintai

| a number of time bases. The most fundamental of these is that of a micro-sec

| from which most other timings are derived. This time base is generated with

| counter clocked at the core clock frequency. Since the core frequency is its

10: | variable, the counter needs to be programmed in accordance with the operatin

| core frequency.

| The TSF is a 64 bits micro-second timer kept by the local station. The TSFs

| all the station in the network are kept roughly synchronised through the

| copying of timestamp carried within the beacon. For an AP, its TSF is the

15: | absolute reference to all other stations and will not be modified once start

| For other stations, it will update its TSF with the timestamp in an error-fr

| beacon unconditionally in a BSS or when the timestamp is later than the TSF

| a IBSS.

20: | There are 3 inter-frame spacings supported by PCU ie. SIFS (16us), DIFS

| (16+18=34us) and EIFS (16+27=43us). These time intervals are measured from t

| moment when the medium has become quiet as indicated by the 'rx_clear' from

| Transmissions are held back until an IFS has elapsed depending what kind of

25: | frame is to be transmitted (ACK, CTS, Fragments use SIFS) or whether the las

| received frame is error free not (use EIFS if error otherwise DIFS). Since I

| includes the transmit latency from MAC to antenna, the IFS's have to shorten

| by the same amount and because the latency is not integral number of

| micro-seconds, the core clock is used to drive the IFS timer to provide

30: | sufficient timing resolution.

| NAV also known as virtual carrier sense is a method to reserve a length of t

| for a station perform its subsequent transmissions with minimum chance of be

| interrupted.

35: | Within the MAC header of every frame type except PS-Poll, there is a 16 bit

| field known as Duration. When a station received a error-free frame not

| directed to it, it will set its NAV to this value if is current NAV is less

| and then decrement it by 1 every micro-second. Whenever its NAV is non-zero,

40: | a station will consider the medium to be busy regardless of what rx_clear

| indicates and refrain from transmitting.

| Random backoff is perform whenever a station attempts to transmit under DCF

| to minimise the chance of collision at the end of a DIFS. A random number

| is chosen within the current contention window and put into the back-off tim

45: | The timer is decremented by 1 every Slot-time (9us) interval and when the ti

| reaches 0, the station can begin its transmission. Since the Slot-time is on

| 9 micro-second long, it is therefore more efficient to derived it from the c

| frequency and is length is controlled by programming the Slot-time timer in

| accordance with the operating core frequency.

50: |

| After the transmission of a uni-cast non-control frame and a RTS frame, the

receiver is expected to respond with an ACK and CTS frame respectively. If the expected response was not forthcoming within a Time-out period, a reception error is assumed. The Time-out periods are expected to be slightly longer than a SIFS and is therefore more efficient for their timer to base off the core clock.

Error Recovery/Transmit Block Mechanism

All successful reception of a uni-cast non-control frame are to be acknowledged by sending ACK frames back to the transmitter. If no expected ACK frame was detected within a ACK Time-out period, the transmitter will schedule a re-transmission of the frame after performing an appropriate back-off. Re-transmissions continue until an ACK is received or the retry limit is reached. In the latter case, the frame is reported as completed but with a FAIL completion status. In addition, the bit corresponding to the receiver Key Index in the Transmit Block register is set. This has the effect of causing all subsequent frames destined to the same receiver to be discarded. These discarded frames are reported as completed with a DISCARD completion status. The block bit can be cleared either by the appearance of a frame with the same Key Index and the ClearDstMask bit of its descriptor set or by writing to the register directly. This mechanism prevents frames from reaching the destination out-of-order and conserves bandwidth by not sending frames to stations that are currently not reachable.

To maximise the probability of success while retrying, the PCU will toggle the antenna select on the 1st retry and then change to the next lower supported rate on the 2nd. This procedure is then repeated for all subsequent retries until success or retry limit is reached.

Fragments

PCU supports fragment burst i.e. successive frames that have the more_frag field set except for the last one. These frames are transmitted in rapid succession interspersed with the ACK frame from the receiver with only a SIFS spacing between them and no random back-off is performed until the frame exchange for the very last fragment is completed. However, if the transmission of any fragment failed, the station would perform back-off and then retry the failed fragment. If successful, the remaining fragments would be transmitted by resuming the SIFS transmission protocol described above. However, if retry limit was reached still without success, then all subsequent fragments of the MSDU must be discarded along with the current one. The Transmit Block mechanism will ensure this is so. Once a fragment burst has begun, it will not be interrupted by other frames even beacons. However, once interrupted because of transmission failure, other higher priority frames can be transmitted before retry.

RTS & CTS

RTS and CTS is supported completely in PCU. On transmit, SW can specify any frame or fragment burst is to be preceded by a RTS/CTS exchange by setting the RTS enable bit in the Transmit Descriptor. PCU will transmit an RTS frame with the Duration field copied from the RTS Duration field in the descriptor and wait for the receiver to respond with a CTS frame. When the CTS frame was received free of error, the Data frame(s) will be transmitted with after a SIFS interval. If the expected CTS did not materialise within a period of CTS-TimeOut, a re-transmission would be scheduled after performing an appropriate back-off. If the RTS/CTS exchange was successful but the Data frame failed, a fresh RTS/CTS exchange would first take place before the frame is re-transmitted.

Beacon and Timers

Beacon is a special management type frame transmitted at 'regular' interval by AP of a BSS or by all stations in an IBSS. It advertises the existence of the network and carries all the essential parameters of the network.

In a BSS, the AP will schedule a beacon for transmission when TBTT (Target Beacon Transmission Time) is reached. TBTTs are separated by Beacon Interval which is a network wide parameter. However the moment of actual transmission must still be dictated by DCF rules.

The content of the beacon in a BSS changes every interval as it contains some fields that count down and fields that indicate buffered traffic for station operating in power saving mode. SW is expected to update the beacon ahead of every TBTT and place it at the head of the Alternate Queue. A timer (SW Beacon Alert) generates an interrupt at a programmable time interval before TBTT to prompt SW to prepare the beacon and the Alternate Queue. Another timer (DMA Beacon Alert) signals the DMA controller at a later time but before TBTT to flush the TX FIFO and begin loading FIFO with frames in the Alternate Queue. After DMA has downloaded the beacon, it will signal to PCU and at TBTT whose arrival is determined by a 3rd timer the beacon will be scheduled. However, if a valid beacon has not yet been loaded at TBTT, then the PCU will wait until a valid beacon signal is detected.

Unlike its counterpart in a BSS, the content of beacons in an IBSS is unchanged once the network is established. Since no AP is present, all stations in the network share the responsibility to transmit the beacon. At TBTT, all stations will contend to transmit the beacon using a modified back-off rule. Whenever a beacon is transmitted, all stations will cancel their scheduled beacon. The station that transmitted the beacon must stay awake throughout the beacon interval to answer Probe Requests.

One of the fields in the Beacon frame body is the Timestamp. At the moment of transmitting the beacon, PCU will insert the current local TSF into it.

PCU supports the following timers which will be used for different purposes according to the role of the station.

Timer0 (1.024 milli-second timer, 16 bits):
 AP & IBSS: Use for generating the TBTT events. The event time is incremented by a Beacon Interval at TBTT to the next TBTT. For an AP or the station that initiates an IBSS, SW should initialise timer to 0 at the moment when the network becomes operational. For a station that joins an IBSS, SW should listen for beacons, and use the knowledge of the beacon interval and the beacon timestamp to compute the next TBTT and initialise the timer with it.

Timer1 (128 micro-second timer, 19 bits):
 AP & IBSS: Use for generating the SW Beacon Alert events. The event time is incremented by a Beacon Interval to the next event time at the current event.
 STA in BSS: Use for generating the wake up from Doze state event. It is incremented by the Listening interval.

Timer2 (128 micro-second timer, 25 bits):
 AP & IBSS: Use for generating the DMA Beacon Alert events. The event time is incremented by a Beacon Interval to the next event time at the current event.
 STA in BSS: Use for generating the Start of Contention-Free Period events. The event time is incremented by a Contention-Free Period to the next event time at the current event.

These timers are enabled simultaneously when the Beacon Interval register is written to.

Flushing TX FIFO

Under some circumstances, the DMA would request PCU to flush the TX FIFO. The PCU will perform the flush immediately except:

1. it is in the middle of an active transmission. This includes waiting for an ACK or CTS and sending a fragment burst.
2. it is in the middle of a retry. The retry will be cancelled and the frame is reported as completed with a retry-limit-exceeded status.

Once the flush has occurred, the flush request will be acknowledged.

Transmission Priority

190: | The 802.11 protocol implies the following transmission priority for differen
| frame types:

195: | 1. ACK/CTS
| 2. Beacon
| 3. PS-Poll
| 4. Other Data/Management Frames

200: | WEP

205: | The WEP unit, shared to encrypt transmitting frames and to decrypt received
| ones, implements the RC4 algorithm with 40 bits or 104 bits key. In addition
| the unit contains a 64 entries key cache which stores keys that are shared
| between the station with other stations. Each entry in the cache holds
| the key, the 48 bits address of the station with which the key is associated
| a bit to indicate the entry is valid or not and a bit to indicate the key
| to be of length 40 bits or 104 bits. Entry 0-3 are reserved for the standard
| 802.11 default keys.

210: | In the transmit direction, PCU will examine the WEP bit in the MAC header.
| If set, the appropriate key is read out from the key cache using the KeyInde
| field in the descriptor as the index. The WEP unit encrypts the frame on the
| fly as the frame is being transmitted. The unit requires 256 cycles for
215: | initialisation and has a throughput of 1 byte per 3 cycles.

220: | Likewise, the PCU examines the WEP bit of a receiving frame to check
| for encryption. If set, the KeyID field of the WEP extension is read.
| If KeyID is non-zero, the default key is looked up from the key cache
| using KeyID as the index. Otherwise, the key cache is searched sequentially
| to match the Transmitter Address with the address in each cache entry.
| If a match is found and the entry was marked valid, the key associated
| with the entry will be used to decrypt the frame body. If no match was
225: | found, then no decryption is performed and this is indicated back to
| SW through the Receive Completion Status. In addition, SW can optionally
| disable the sending of ACK responses for encrypted frames for which no
| key can be found.

230: | Receive Filter

235: | PCU supports 5 types of receive filter. They are:

240: | 1. Unicast
| When enabled, Data and Management frames with a Receiver Address matching
| the station address would be passed to SW.

245: | 2. Multicast
| When enabled, Data and Management frames that have a Multi-cast Address
| (bit 47 set) which passes a hashing test would be passed to SW.
| Hashing test:
| The 48 bits address is divided in to 8 6 bits segments and then
| bit-wise XOR all 8 segments to arrive at a 6 bit number. Use this
| to look up the corresponding bit in the 64 bits Multi-cast Filter
| Vector which is programmable by SW. If the bit is set, the test
250: | is considered passed.

255: | 3. Broadcast
| When enabled, Data and Management frames with a Receiver Address of all
| 0xffffffff would be passed to SW.

255: | 4. Control
| When enabled, all Control type frames detected in the medium would be
| passed to SW.

255: | 5. Promiscuous
| When enabled, all frames detected in the medium would be passed to SW.

| In addition, any Beacon frames regardless of BSSID and unicast PS-Poll

frames are passed to SW un-filtered.

260: |

| TIM & PSpoll

265: | In a BSS, the AP buffers any frames that are destined to a Power-Save capable station which is in Doze state. The station will come out of dozing every Listen Interval to listen for beacon from the AP. Carried in the beacon frame body is a TIM element with which the AP indicates the presence of traffic for all the associated PS station. In the event where TIM indicates the presence of buffered frames, the station is expected to

270: | respond with sending a PS-Poll frame to the AP prompting it to transmit the frames in its buffer.

| To facilitate the parsing of the TIM element by PCU, the offset in bytes from the start of the MAC header to the Bitmap Control sub-field in the TIM

275: | will be programmed by SW.

| Co-existence with PCF

280: | At the beginning of a Contention Free Period (CFP), NAV will be set to CFPMaxDuration, a constant which can be found in the CF parameter element of a beacon. When a CF-End/CF-End-Ack frame is detected, NAV will be reset to zero. This means the intermediate updates of NAV with CFPDurRemaini will not be performed in our implementation.

285: | Ad hoc

| Before TBTT arrives, a PS station will be wakened by the DMA beacon alert in preparation for sending and listening for beacons and ATIMs.

290: | At TBTT, a beacon transmission is scheduled using a contention window of $2 \times CW_{min}$ (don't know what should happen if the delay chosen was 0). During the back-off count-down, if a beacon was received then the scheduled beacon would be cancelled. Note that a fragment burst will not be interrupted in order to send a beacon.

295: | During the ATIM window which last for an ATIM window interval from TBTT, no transmission of frames other than beacon, ATIM, ACK, RTS & CTS will be initiated but a frame started before TBTT is allowed to complete. ATIMs (if any) are loaded into the Tx FIFO after DMA beacon alert. After either

300: | a beacon is sent or received, ATIM frames are read out from the FIFO and sent. When the first non-ATIM frame emerged from FIFO, then transmission was suspended until the ATIM window has ended. On the other hand, any ATIMs left in the FIFO after the end of ATIM window will be discarded.

305: | As part of the preparation for the up-coming beacon period, SW will set bits of the Tx Destination Block register that correspond to all PS stations in the network. Only an acknowledged transmission of ATIM directed that destination can clear the block allowing subsequent frames for that destination to be sent. For broadcast or multicast traffic, an key index and

310: | hence a Tx block bit must be allocated to each session. Clearing the bit in this case only requires the transmission of the corresponding ATIM.

| When a station has received a relevant ATIM or transmitted a beacon, it would remain awake until the end of the next ATIM window. Otherwise, it will go

315: | into power save mode. If there are frames to be transmitted, SW would need to wake the station manually.

|

| Support for Power Save

320: | TBD

|

| Features to support Diagnostic & MIB

325: | The following features are not functional but purely diagnostic in purpose:

| 1. Disable ACK response


```

330:      |2. Disable CTS response
      |3. Disable Encryption
      |4. Disable Decryption
      |5. Disable Transmit
      |6. Disable Receive
      |7. Loop back
      |
335:      |The following information is accumulated in PCU and is accessible via
      |register read:
      |
      |1. RTS success count
      |2. RTS failure count
340:      |3. ACK failure count
      |4. FCS check failure count
      |5. Current Timestamp
      |6. Timestamp in the last received beacon
      |
345:      |
      |Outstanding Issues:
      |
      |- Can SW send control packets
      |- Power Save Mode
350:      |
      |
      |
355:      |
      |
      |

```

P4 admin: Jeff Smith

1. A Key Caching Mechanism

In 802.11 a key may be used to encrypt packets. A unique key may be assigned per transmit-receive combination. An Access point (AP) connects to many Stations (STA) so therefore may be required to store many keys. The AP must use the STA address that the packet is to be transmitted to or has been received from to determine which key to use. It must search through all available address-key combinations to find the correct element. Since keys may be 128 bits or greater, the addresses 48 bits and an AP may be required to support the connection of many STAs this table of key-address tuples may grow very large. Yet, it is desirable to keep the size small as it lessens the maximum amount of time required to search the table and minimizes the cost of the solution.

This scheme maintains a table of *active* key-address tuples which can be searched as required. It also keeps a table of *stored* key-address tuples which can be activated when required.

Many communication systems require a positive acknowledgement of data that has been sent. In such a system data is sent to the destination which must generate an acknowledgement if it receives the data successfully. If no such acknowledgement is received the source will resend the data some number of times.

This scheme takes advantage of such a positive acknowledgement based system by NOT returning the acknowledgement for received data if the source address was not found in the active table. This gives time for the receiving system to examine its stored table to see if the required tuple exists there. If not, the system may either continue to ignore the data, by not returning an acknowledgement, or program a fake tuple into the active table so that an acknowledgement is returned before discarding the data. If the source address is in the stored table then it may be moved to the active table to receive the resent data. The key should then be found successfully and the acknowledgement generated.

This scheme applies to IEEE802.11 since it uses keys to encrypt/decrypt packets and requires receiving STAs to generate and acknowledgement packet after successfully receiving a packet.